Machine Reading as a Cognitive Science Research Instrument

Kenneth D. Forbus, Kate Lockwood, Emmett Tomai, Morteza Dehghani, Jakub Czyz

Qualitative Reasoning Group, Northwestern University
2133 Sheridan Road, Evanston, IL, 60208, USA
{forbus, kate, etomai, morteza, jake}@northwestern.edu;

Abstract

We describe how we are using natural language techniques to develop systems that can automatically encode a range of input materials for cognitive simulations. We start by summarizing this type of problem, and the components we are using. We then describe three projects that are using this common infrastructure: learning from multimodal materials, modeling decision making in moral dilemmas, and modeling conceptual change in development.

Introduction

Many of today's cognitive simulations use inputs that are hand-encoded, sometimes by the experimenters themselves. This leads to two significant problems. The first is that it increases the possibilities for *tailorability*, where the results of a computational experiment derive more from free parameters in the experiment than on the hypotheses being tested. The second is that it limits the scale of simulation experiments which can be undertaken. For example, studying how decision-making is influenced by cultural background or studying how conceptual change occurs in development requires formally encoding a large body of material, a daunting prospect that has greatly restricted the scope of simulation experiments. In this paper we describe how we are using natural language techniques to develop systems that can automatically encode a range of input materials for cognitive simulations. We start by summarizing this type of problem, and the components we are using. We then describe three projects that are using this common infrastructure: learning from multimodal materials, modeling decision making in moral dilemmas, and modeling conceptual change in development. We believe each of these projects, aside from being interesting uses of machine reading, have potential for generating results that will be useful for learning by reading, and thus will be of particular interest to participants in this symposium.

The Task: Understanding texts in cognitive science experiments

Cognitive simulation experiments attempt to model some aspect of human behavior. The particular experimental tasks used vary widely, as do the measures used to evaluate results. However, there are two common kinds of text that appear in experimental situations: (1) *Input stimuli* given to participants in experiments often takes the form of single sentences, paragraphs, or even short stories/articles, sometimes combined with diagrams. Ideally, one would like to give both people and programs the identical input materials, and carry out the analogous measures on each, in testing a theory. Typically, however, inputs are hand-translated into formal representations, which are then given to simulations. (2) *Interviews* with participants are typically transcribed into text, and provide more detailed windows into what people know and the explanations that they give for their conclusions. For example, in interviews with groups from different cultures, one might ask how a change in a population within an ecosystem will affect other populations. Analyzing such protocols is currently typically done by hand, annotating by some coding scheme segments that are viewed as being significant relative to the hypotheses being tested.

A third kind of text found in cognitive simulation experiments encodes *background knowledge/experience* for the simulation. This is rare in practice; one of the few examples of which we are aware is John Anderson's use of a controlled language to express simple procedures, which are automatically translated into an internal declarative representation for their ACT-R skill learning simulations. Obviously, people have a great deal of experience, relative to any existing cognitive simulation system. Being able to provide more simulated background via simple natural language texts would enable the simulation of larger-scale learning phenomena, such as misconceptions arising from learning and conceptual change.

Our Approach

As noted above, the workflow today for all three types of text is essentially manual labor. This is a serious bottleneck. We believe that this bottleneck can be eliminated by using extensions to off-the-shelf natural language technology and controlled languages. Let us

Report Documentation Page

1. REPORT DATE **2007**	2. REPORT TYPE	3. DATES COVERED **00-00-2007 to 00-00-2007**

4. TITLE AND SUBTITLE **Machine Reading as a Cognitive Science Research Instrument**	5a. CONTRACT NUMBER
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Northwestern University,Qualitative Reasoning Group,2133 Sheridan Road,Evanston,IL,60208**	8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT **Approved for public release; distribution unlimited**

13. SUPPLEMENTARY NOTES

14. ABSTRACT
We describe how we are using natural language techniques to develop systems that can automatically encode a range of input materials for cognitive simulations. We start by summarizing this type of problem, and the components we are using. We then describe three projects that are using this common infrastructure: learning from multimodal materials, modeling decision making in moral dilemmas, and modeling conceptual change in development.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT **Same as Report (SAR)**	18. NUMBER OF PAGES **6**	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**			

examine why, for each kinds of text in turn. (1) Input stimuli in most experiments is kept fairly simple, so that reading difficulties do not confound the experimental results. The texts in experiments that our collaborators are carrying out tend to be simple enough that they fit within our controlled language, or can easily be re-written to do so. (2) Interviews provide information about the models and experience of a particular participant. They are often ungrammatical, and the transcription process can introduce errors. Thus a semi-automatic tool (described below) that helps experimenters translate them into our controlled language is important. (3) Background texts, since they are never seen by human participants, can be written as needed to fit our controlled language conventions. As long as the cost of writing simplified text is less than the cost of encoding knowledge in predicate calculus by hand, automatic processing will be worthwhile.

Next we briefly review the components that we are building on in all three of the projects described below.

Explanation Agent NLU System: This system was originally constructed by Kuehne *et al* (2002, 2004) as part of an exploration of how qualitative process theory (Forbus 1984) could be used in natural language semantics. Kuehne showed that there are mappings between linguistic constructions involving quantities and causal relationships between them that are naturally captured by QP theory, demonstrating the utility of his ideas via a corpus analysis and constructing a controlled language NL system that could produce QP descriptions from texts. EANLU uses Allen's parser (Allen, 1995), the COMLEX lexicon, a knowledge base derived from the contents of ResearchCyc, and the FIRE reasoning engine developed at Northwestern. The controlled language it handles is QRG-CE, for "QRG Controlled English". Unlike controlled languages used for authoring and translating manuals, such as Caterpillar English (Kamprath *et al* 1998), QRG-CE allows the introduction of new vocabulary terms, to support learning experiments. We are currently extending it to handle the range of contents described below.

SME: The Structure-Mapping Engine (Falkenhainer, Forbus, & Gentner, 1989) provides a model of analogical matching, based on Gentner's (1983) structure-mapping theory. SME has been used to model a variety of psychological phenomena, and has been used to make novel predictions that subsequently have been borne out in laboratory experiments.

MAC/FAC: This model of similarity-based retrieval (Forbus, Gentner, & Law, 1995) uses a first stage that operates on simple feature vectors, automatically constructed from structured representations, and a second stage that uses SME, to automatically retrieve situations that are similar to whatever a system is reasoning about. MAC/FAC has been used to model several phenomena in the psychology of remindings.

SEQL: This model of generalization (Kuehne, *et al* 2000) uses SME to construct generalizations from examples. It has been used to model psychological phenomena involving sequence effects in learning, and is capable of learning at rates that are more consistent with human learning than today's purely statistical learning systems. Recently SEQL has been extended with probabilities (Halstead & Forbus, 2005), as a means of combining analogical and probabilistic inference.

sKEA: The sketching Knowledge Entry Associate (Forbus & Usher, 2001) is the first open-domain sketch understanding system. It uses the same ResearchCyc-derived knowledge base, and enables users to conceptually label their ink with its intended meaning. Its visual processing is crafted to be consistent with what is currently known about human high-level visual perception.

Using Multimodal Communication for Knowledge Capture

Many textbooks use a combination of diagrams/drawings and text to help convey ideas and information, especially in science and engineering. Consequently, being able to integrate these sources of information is an important problem in knowledge capture. For example, consider the following selection from *Sun Up to Sun Down* (Buckley, 1979), an introductory text on solar energy:


Figure 1. Text + diagram example from *Sun Up to Sun Down*

While there are many deep problems involved in understanding such diagram/text combinations (e.g., Novak & Bulko, 1990; Hegarty & Just, 1993), one central problem is correctly understanding spatial prepositions. Spatial prepositions convey relationships between objects. In a multimodal context (i.e., sketch plus language), spatial prepositions must be correctly understood to disambiguate references to items in a sketch and to help create a model for the sketched objects.

Background. Utterances involving spatial prepositions involve at minimum two objects: a reference object (the ground) and a located object (the figure) as well as the preposition that describes their relationship. Many recent psychological studies have focused on understanding which properties of the figure and ground objects play a role in the assignment of these prepositions (Feist & Gentner, 1998; Carlson-Radvansky et al, 1999).

Learning Spatial Language. Lockwood (*et al* 2006) has shown that some spatial prepositions can be automatically

learned from sketches using SEQL. In these experiments we took unlabeled sketches of simple geometric shapes representing five different prepositions: *in, on, above, below* and *left* and processed the representations created to extract the qualitative features necessary to create categories corresponding to the prepositions portrayed. All of the stimuli used were drawn from psychological experiments where people or other programs were learning and labeling spatial relationships.

Our next step is to run the same experiment with sketches portraying real-world objects in functional situations, for example: an apple in a bowl. Functional features as well as topological ones have been shown to play an important role in spatial preposition use (Coventry & Garrod, 2004). We also plan to run experiments to see if we can learn spatial prepositions from another language, since systems of spatial prepositions vary significantly between languages (Bowerman, 1993).

Using Spatial Language. Results from the learning experiments will be used to inform SpaceCase (Lockwood, et al, 2005), a model that uses Bayesian rules to identify spatial prepositions in a sketch. SpaceCase has been used to correctly label sketches that combine functional and geometric information, including memory effects imposed by linguistically labeling geometrically ambiguous stimuli (Feist & Gentner, 1998). Previously SpaceCase used rules collected from the psychological literature. Next we plan to experiment with generating the rules from the classifications we create in the learning experiments.

We are implementing the multimodal interaction system in our Companions cognitive systems architecture (Forbus & Hinrichs, 2006). Our sKEA sketch understanding system is already integrated into Companions. The EA NLU system is being packaged up as another agent on the cluster, used as a service by the Interaction Manager. To handle these inputs, the Interaction Manager is being extended with a TRIPS-style dialogue manager. A chat interface on the client machine, connected to the Session Manager, provides textual input/output capabilities. We plan to use the EA NLU system both to parse full texts/stories for knowledge capture and to parse and respond to user utterances and questions.
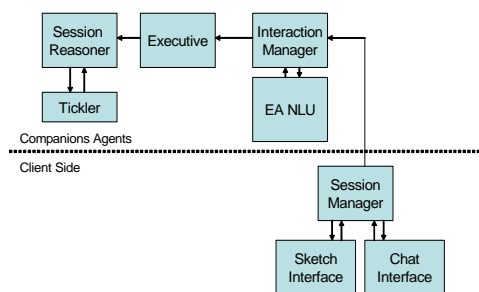

Figure 2. Companions architecture, with NL extensions.

In addition to giving our natural language system greater ability to understand texts, a working knowledge of spatial prepositions will also allow it to query the user to clarify ambiguous or contradictory statements during multi-modal knowledge capture. For example in Figure 1 above, if the system is told "the water level drops" this is a potentially confusing piece of information – there are several instances of water in the diagram. The system could then query the user "which water are you talking about?" and display the text in question. The user could respond "the water *in the* pipe" using a spatial preposition to specify the important entity. Now the system knows that the "water level *in the pipe* drops" making its understanding more complete.

Our model for the workflow with the system is as follows. We provide a coherent piece of text, written in QRG-CE, with one or more diagrams. The system processes this material, and is allowed to ask questions about ambiguities, either due to its understanding of the lesson contents, or problems with understanding the text and/or the diagrams, individually or in concert. We then ask the system further questions, if desired. The material from the lesson will be stored in the system's KB, to be used in subsequent learning. We also plan to experiment with rumination processes (cf. Forbus *et al* 2007).

To test the system's ability to capture knowledge, we plan to feed it the complete text from *Sun Up to Sun Down* rewritten into QRG-CE, along with sketches of all of the diagrams in the book. (This book contains over 1,200 sentences, divided into 25 chapters, with one diagram per page, on average. Translating to QRG-CE will more than double the number of sentences.) To ensure robustness, we also plan to do the same for a middle-school science textbook on heat energy. We have obtained the teacher's manual and quizzes for this book, which will provide an external benchmark for evaluating the system's understanding gleaned by reading.

Modeling Decision-Making in Moral Dilemmas

In joint work with psychologist Doug Medin's group at Northwestern University, we are exploring the cognitive processes underlying moral reasoning. Medin's group has been working for some time with several populations, including three distinct cultural groups in Guatemala, and with Amish, Menominee Indians, and majority culture people in Wisconsin. They are investigating their mental models of nature, used in decision-making (cf. Medin & Atran, 2004), and how people reason about moral dilemmas. Our goal is to create computational models of such decision-making. We will be using semi-automated reading of text from interviews done in conjunction with survey questions about morally difficult choices. These surveys are being done across various cultural groups to investigate the impact of cultural values and influences on the reasoning process (Tanner, Medin & Iliev, forthcoming).

Traditionally, research on human decision making has mainly been focused on secular goods. More recently, the concepts of *protected values* (Baron & Spranca, 1997) or *sacred values* (Tetlock *et al* 2000), have been developed in

Forbus, K., Lockwood, K., Tomai, E., Dehghani, M. and Czyz, J. (2007). Machine Reading as a Cognitive Science Research Instrument. AAAI Spring Symposium on Machine Reading. Stanford University, California.

accordance with the new evidence concerning the moral aspect of choice. Unlike secular goods, these values are usually seen as absolute and non-changing. These moral values influence people's information processing and decision making. As a result of these influences two different perspectives on decision making emerge: *deontological* and *consequentialist*. Deontological decisions are made based one's duties or rules and the rights of others. Duties are defined as morally mandatory actions or precipitations, such as the duty not to kill. In contrast, in consequentialist decision making the emphasis is on the consequences of actions and not on the nature of it. Therefore, the conclusion about whether a certain action is good or evil is determined by its consequences and overall utility. For example, if 100 children have been diagnosed with a certain disease and there is a vaccine which would cure the disease, but also it will kill 1% of the recipients of the vaccine, a deontological decision maker would not agree with children being vaccined. A consequentialist decision maker on the other hand, would calculate the overall utility of the action, and given that vaccination of children leads to a higher overall utility, would agree with the decision. These can lead to strong individual differences; for example, reasoning from strong protected values can lead to immunity to framing effects that can otherwise distort decision-making (Tanner & Medin, 2004).

Experimental goals and approach. We plan to use the EA NLU system to read controlled language translations of data from interviews. This data will provide two kinds of information. First, in some cases it will provide direct evidence for qualitative models held by the participants (e.g., "the more pollution generated, the more fish will die"). These qualitative models will be used in the simulation to predict answers and explanations given in response to novel questions and dilemmas. Second, it will provide a form of data about background experience, including in some cases the stories that a culture refers to as analogs in decision-making. These also will be used in the simulation, retrieved via MAC/FAC and applied via SME to make predictions about judgments in new situations. The models will be formed for individuals but also generalized over sets of individuals belonging to the same cultural group, using SEQL.

Interviews and controlled language. The interview transcripts provide a wealth of insight into the reasoning processes of the participants, but encoding them in a formal knowledge representation is a very difficult task. Participants in interviews often speak ungrammatically, using fragments and backtracking to fill out an explanation. The range of topics that come up in the moral dilemmas and ecosystem questions are also quite broad. These factors suggest that fully automatic encoding is unlikely to be successful. Consequently, we are using a semi-automatic approach. Previous work with the EA NLU system demonstrated the effectiveness of using controlled language for building qualitative models of physical processes from texts. Translating from the unconstrained, wide-open text in interviews to a controlled language, especially using a workbench that provides ample feedback about the process, seems to us to be a practical compromise. Our goal is to make it easier for experimenters to produce formal representations from interviews compared to doing it by hand, a process which is well-known to be extremely laborious and error-prone.

Extending our representations. Fortunately, the contents of the ResearchCyc knowledge base we are using gives us a broad set of everyday knowledge beyond the QP-style knowledge that the system already handles. Looking over the interview data gathered so far, it appears that many of the concepts we will need are already in the KB, although some extensions will be necessary. The ResearchCyc vocabulary for spatial, geographical, and temporal knowledge, for example, seems to be adequate for our needs, and many of the event types we need are also already there. However, there are several areas where substantial knowledge extensions appear to be necessary. We plan to use the EMA model of emotions (Gratch *et al* 2006) to help model value judgments and their affective consequences. Participants often use stories to illustrate their points, so we are creating representations for narrative elements such as characters, scenes, plot progression, paralleling, climax, and morals.

Extending QRG-CE and EA NLU. This project requires substantial extensions to the grammar of our controlled language and semantic interpretation processes. After all, in physical scenarios, there are not agents with desires, goals, beliefs, and capable of making proclamations. Our approach is minimalist and practical: We are making the smallest extensions we can, consistent with covering the kinds of materials we are finding in interviews. (We are also examining fables and short stories drawn from a diverse set of cultures, to ensure that our system does not become too specialized.)

Supporting the translation process. We are working in collaboration with Mark Finlayson and Patrick Winston at MIT to build a workbench to support the translation of interview data into formal representations, using QRG-CE as an intermediate language. This is inspired by both CMU's KANT project (cf. Nyberg *et al* 2002) and Boeing's controlled language work (cf. Clark *et al* 2003). Our goal is to provide copious feedback to aid researchers in the translation process, by (1) making the user immediately aware of errors, and (2) making it clear how they might change the input text to resolve the problem. For each sentence of input, the system will display the entities and events referenced (explicitly or implicitly) along with role assignments and spatial, temporal and causal relationships that it believes are described by the sentence. Instances of missing relationships can be highlighted where a role relation is unfilled or an event lacks spatial, temporal or causal connections with other events. Ambiguous interpretations resulting in multiple possible models will be presented for manual disambiguation when necessary. EA NLU already has an interactive mode where feedback about ambiguities is used

Forbus, K., Lockwood, K., Tomai, E., Dehghani, M. and Czyz, J. (2007). Machine Reading as a Cognitive Science Research Instrument. AAAI Spring Symposium on Machine Reading. Stanford University, California.

for evidential reasoning about potential future ambiguities inspired by (Barker *et al* 1998), and manual disambiguation data will be used for further training. For testing the effectiveness of the workbench, we will use newly gathered transcripts, since experiments are on-going. As we bootstrap our cognitive simulations with material from existing interviews, we will use them to make predictions which will help guide subsequent data-collection efforts.

Modeling Conceptual Change in Development

The causal models children have change considerably as they experience the world and interact with others. Developmentally, children often appear to initially use simple essence models. For instance, children at 4-5 years often believe that things float because they are stronger than water. They believe that this one central property is predictive of a wide range of possible behaviors and effects. At this stage, children tend to predict that nails will float because they are "stronger" than water. They are astonished when they place a nail in water and watch it sink. Eventually they learn otherwise, but it is not always clear what is learned: Even many adults find the idea of concrete boats implausible, unless they explicitly think about steel boats (e.g., freighters). We think that comparison plays a vital role in reshaping causal representations to become more subtle and articulate, by helping children (and adults) explain things to themselves. Conservation phenomena provide another class of examples: Children learn at different ages that (a) the number of objects doesn't change if they are piled up versus distributed, (b) that a lump of clay is the same amount of stuff if it is rolled up versus spread out, and (c) that the volume of a liquid is preserved over pouring it from one container to another, even if the level changes due to the shape of the container.

Example-based learning is arguably the most effective natural way to learn a domain. But it comes with an initial cost. Novice learners, both children and adults, typically have highly specific, contextually situated models of phenomena. They may produce totally different explanations for evaporation from a puddle vs. from a clothesline, for example (Collins & Gentner, 1987). At this stage, the knowledge of important principles may be essentially inert: that is, knowledge may be encoded in a domain-specific manner, such that the learner is unlikely to be reminded of it except in case of very strong surface similarity. One way to overcome this initially conservative learning is through comparison with further exemplars, as demonstrated in (Gentner, Loewenstein & Thompson, 2003), which examined business school students learning negotiation strategies. When students compared two cases, they were able to use the material more effectively. Our hypothesis is that the comparison process forced students to re-represent the cases in ways that facilitated subsequent matching, a process we call *analogical encoding*.

By its very nature, the ability to do automatic encoding of stimuli is essential for exploring analogical encoding. Our goal in this effort is to model processes of cognitive development, by using combinations of simplified text and sketches to simulate the experiences that a learner has with particular physical phenomena. For example, in the floating experiments described above, a comic strip-like sequence of sketches will be used to illustrate the before and after situations when the nail is released on the surface of the water. The sketches will provide the spatial relationships, and text will be used to provide other, non-spatial information, such as "The nail is made of iron. The nail is heavy." The puddle/clothesline example illustrates the need to draw upon a broad range of everyday concepts.

Of particular interest is the role of explanations in learning. What makes learners satisfied, or dissatisfied, with their own explanations? What criteria should be used in changing one's explanatory theories? Modeling such phenomena is complex, since it requires looking at larger grain-size phenomena than most cognitive simulation efforts. It also requires the ability to generate a large body of complex stimuli, to serve as a model of experience. The initial descriptions need to be in concrete, particular terms. We are extending sKEA and EA NLU to generate a set of stimuli for the simulation. The automatic construction of the underlying formal representations is a useful way to reduce tailorability, and to explore alternate hypotheses about human encoding processes. We plan to experiment with different strategies for constructing, accepting, and rejecting explanations, using different sequences of phenomena and human data to find potentially plausible mechanisms.

Discussion

We have summarized three projects in progress, all of which are using the same substrate of natural language, sketch understanding, and analogical processing techniques to model complex cognitive phenomena. While these efforts are very much in progress, we think they are illustrative of the kind of project that can be undertaken, given recent advances in the state of the art. Machine reading, we believe, could turn out to be a source of extremely valuable instruments for cognitive science research.

Acknowledgements

References

Barker, K., Deslisle, S., and Szpakowicz, S. 1998. Test-driving TANKA: Evaluating a semi-automatic system of text analysis for

Forbus, K., Lockwood, K., Tomai, E., Dehghani, M. and Czyz, J. (2007). Machine Reading as a Cognitive Science Research Instrument. AAAI Spring Symposium on Machine Reading. Stanford University, California.

knowledge acquisition. *Proceedings of the 12ᵗʰ Canadian Conference on Artificial Intelligence.* pp 60-71.

Baron, J. & Spranca, M. 1997.Protected values. Organizational Behavior and Human Decision Processes, 70, 1-16.

Buckley, S. 1979. *Sun up to sun down.* New York: McGraw-Hill.

Bowerman, M.. (1996). Learning How to Structure Space for Language: A Crosslinguistic Perspective. In P. Bloom, M.A. Peterson, L. Nadel, & M.F. Garrett (eds.) *Language and Space* (493-530). Cambridge, Mass. MIT Press

Carlson-Radvansky, L.A., Covey, E.S., & Lattanzi, K.M. (1999). "What" effects on "Where": Functional influences on spatial relations. *Psychological Science*, 10(6): 516-521.

Clark, P., Harrison, P. and Thompson, J. 2003. A Knowledge-Driven Approach to Text Meaning Processing. *Proceedings of the HLT Workshop on Text Meaning Processing*, pp 1-6.

Collins, A., & Gentner, D. (1987). How people construct mental models. In D. Holland & N. Quinn (Eds.), *Cultural models in language and thought* (pp. 243-265). Cambridge, England: Cambridge University Press.

Coventry, K.R., & Garrod, S.C. (2004). *Saying, Seeing, and Acting: The Psychological Semantics of Spatial Prepositions.* Lawrence Erlbaum Associates.

Feist, M.I., & Gentner, D. (1998). On Plates, Bowls, and Dishes: Factors in the Use of English 'in' and 'on'. *Proceedings of the 20ᵗʰ Annual Conference of the Cognitive Science Society*.

Forbus, K. 1984. Qualitative process theory. Artificial Intelligence, 24, 1984.

Forbus, K., Gentner, D. & Law, K. 1995. MAC/FAC: A model of Similarity-based Retrieval. Cognitive Science, 19(2), April-June, pp 141-205.

Forbus, K. and Hinrichs, T. 2006. Companion Cognitive Systems: A Step toward Human-Level AI. *AI Magazine*, 27(2): Summer 2006: 83-95.

Forbus, K., Riesbeck, C., Birnbaum, L., Livingston, K., Sharma, A., Ureel, L. 2007. A Prototype System that Learns by Reading Simplified Texts. AAAI Spring Symposium on Machine Reading.

Forbus, K. and Usher, J. (2002). Sketching for knowledge capture: A progress report. *IUI'02*, January 13-16, 2002, San Francisco, California.

Gratch, J., Marsella, S., and Mao W. 2006. Towards a Validated Model of "Emotional Intelligence." *Twenty-First National Conference on Artificial Intelligence (AAAI06)*.

Gentner, D. 1983. Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155-170.

Gentner, D., Loewenstein, J., and Thompson, L. (2003). Learning and transfer: A general role for analogical encoding. Journal of Educational Psychology, 95(2), 393-408.

Halstead, D. and Forbus, K. (2005). Transforming between Propositions and Features: Bridging the Gap. *Proceedings of AAAI-2005*. Pittsburgh, PA.

Hegarty, M., and Just, M. 1993. Constructing mental models of machines from text and diagrams. *Journal of Memory and Language* **32**, 717-742**.**

Kuehne, S., Forbus, K., Gentner. D., & Quinn, B. 2000. SEQL: Category learning as progressive abstraction using structure mapping. *Proceedings of the 22ⁿᵈ Annual Meeting of the Cognitive Science Society*.

Kuehne, S.E., and Forbus, K. D. (2002). Qualitative physics as a component in natural language semantics: A progress report. *Proceedings of the 24th Annual Meeting of the Cognitive Science Society*, George Mason University, Fairfax, VA.

Kuehne, S. E. (2004). On the Representation of Physical Quantities in Natural Language Text. *Proceedings of the Twenty-sixth Annual Meeting of the Cognitive Science Society*, Chicago, Illinois, USA, August.

Lockwood, K., Forbus, K., & Usher, J. (2005). SpaceCase: A model of spatial preposition use. *Proceedings of the 27ᵗʰ Annual Conference of the Cognitive Science Society*. Stressa, Italy.

Lockwood, K., Forbus, K., Halstead, D., & Usher, J. (2006). Automatic Categorization of Spatial Prepositions. *Proceedings of the 28ᵗʰ Annual Conference of the Cognitive Science Society.* Vancouver, Canada.

Medin, D.L & Atran, S. (2004). The Native Mind: Biological Categorization, Reasoning and Decision Making in Development Across Cultures. *Psychological Review*, 111(4), 960-983.

Novak, G. and Bulko, W. 1990. Understanding natural language with diagrams. *Proceedings of AAAI-90*

Nyberg, E., Mitamura, T., Baker, K., Svoboda, D., Peterson, B., and Williams, J. 2002. Deriving Semantic Knowledge from Descriptive Texts using an MT System. *Proceedings of AMTA 2002.*

Kamprath, C., Adolphson, E., Mitamura, T. and Nyberg, E. 1998. Controlled Language for Multilingual Document Production: Experience with Caterpillar Technical English. *Proceedings of the 2ⁿᵈ International Workshop on Controlled Language Applications (CLAW'98)*

Tanner, C. and Medin, D.L. 2004. Protected Values: No omission bias and no framing effects. *Psychonomic Bulletin and Review*, 11(1), 185-191.

Tanner, C., Medin D. L., Iliev R. (Forthcoming) "Influence of Deontological vs. Consequentialist Orientations on Act Choices and Framing Effects: When Principles are more Important than Consequences".

Tetlock, P.E., Kristel, O.V., Elson, S.B., Green, M., & Lerner, J.S. (2000). The psychology of the unthinkable. Taboo trade-offs, forbidden base rates, and heretical counterfactuals. Journal of Personality & Social Psychology, 5, 853-870.

Forbus, K., Lockwood, K., Tomai, E., Dehghani, M. and Czyz, J. (2007). Machine Reading as a Cognitive Science Research Instrument. AAAI Spring Symposium on Machine Reading. Stanford University, California.